

METHOD AND APPARATUS FOR EFFICIENT
PROTOCOL-INDEPENDENT TRUNKING OF DATA SIGNALS

CROSS-REFERENCE TO RELATED APPLICATIONS

This is the first application filed for the present
5 invention.

MICROFICHE APPENDIX

Not applicable.

TECHNICAL FIELD

The present invention relates to digital
10 communications networks, and in particular to a method and
apparatus enabling efficient protocol-independent extension
of data services through a broadband packet network.

BACKGROUND OF THE INVENTION

Communication networks are constructed using
15 network models and transfer protocols as guides. Network
models and transfer protocols have proliferated over the
past few years as new networks have been developed (and
existing networks have evolved) to accommodate new
end-user, device and application requirements.

20 The proliferation of network models has focused on
monolithic, highly integrated network and network component
architectures. Data services (e.g. functionality and
products) developed for a network that supports a target
group of end-users, devices or applications are usually
25 available only within the geographical reach of that
network, and are not easily extendible to another network
without extensive modification or complete re-creation.

Adding new functionality to installed networks requires complex modifications or overlays to the existing functionality. Conversely, new networks cannot capitalize on functionality and protocols developed for legacy networks, because monolithic integrated designs do not permit porting without complex protocol conversion interfaces. The time required to implement such modifications is incompatible with the modern competitive communications business environment, where the time to market can significantly influence market share. At the same time, the costs associated with installing new network equipment can pose a significant impediment to the deployment of modern network services. For example, many firms (typically those with multiple branch locations) have deployed enterprise networks linking their various branches via lines leased from a network service provider. Legacy enterprise networks of this type were commonly constructed based on any of the X.25, Frame Relay (FR), or Integrated Services Digital Network (ISDN) protocols, and usually involved substantial investment in networking equipment. This investment may make the firm reluctant to incur the cost of new equipment required to utilize modern broadband packet network architectures (e.g. ATM or XGb Ethernet or Internet Protocol [IP]), in spite of the various advantages offered by the newer architecture to the firm and/or the network service provider.

As a result, the present communications network space is a patchwork of modern and legacy networks, each operating under respective network models and transfer protocols (e.g. Ethernet, Time Division Multiple Access (TDMA), frame relay (FR), synchronous transfer mode (STM),

asynchronous transfer mode (ATM), X.25, SNA, Video, Transfer Control Protocol/Internet Protocol (TCP/IP) etc.); utilizing different transport media (e.g. copper, fiber, wireless etc.); and frequently owned by different business entities. Networks may be edge-connected by means of gateway interfaces, permitting transport of data between the networks. However, where such networks are dissimilar, adaptation is required to facilitate the transport of data across the networks. In the context of the present invention, adjoining networks are considered to be "dissimilar" or "heterogeneous" if their respective network models, transfer protocols and/or media are sufficiently different that data streams originating in one network cannot be transported by an adjoining network without conversion or reformatting. Such conversion, referred to herein as "adaptation", is an automated process to enable the transfer of data across the heterogeneous networks.

For example, a firm may operate a legacy enterprise network comprising a site network at each of its various branch locations, and utilizing frame relay for inter-site communications over leased lines provided by a network service provider. The network service provider may wish to extend the frame relay service through an ATM backbone network, in order to take advantage of the superior data transport capabilities of the ATM backbone. Such service extension requires adaptation of data at gateways between the ATM network and each of the involved site networks. Thus data originating in each site network (and conforming to the FR protocol) must be adapted to conform to the ATM protocol before it can be transported across the ATM backbone. Similarly, data arriving at a site network

through the ATM backbone must be adapted to conform to the FR protocol before the data is delivered to the site network.

Adaptation to facilitate data service extension scenarios of this type is typically performed in gateway servers provided for that purpose, usually as part of a data services solution specifically engineered to the requirements of the customer. Both of these factors tend to increase costs and reduce flexibility, thereby creating an impediment to their widespread deployment.

Applicants co-pending U.S. Patent Application No. 09/158,855, entitled TRANSIT TRUNK SUBNETWORK, which was filed on September 23, 1998, teaches a method of inter-office trunking of PCM signals through an ATM backbone. A PCM data stream received at an ingress gateway is packed into ATM cells having a predetermined length. The ATM cells are transported through the ATM backbone, e.g. over a switched virtual circuit (SVC), to an egress gateway. The egress gateway extracts the data from the ATM cells to recover the original PCM signal.

ATM SVCs are particularly adapted for trunking circuit-switched data streams, because ATM guarantees delivery of each cell, and ensures that cell-ordering is maintained. Thus ATM can readily satisfy the QoS requirements of connection-oriented traffic without excess overhead. However, in some cases it may be desirable to provide trunking through other packet-based networks (such as IP, Ethernet or wireless), in which QoS guarantees may not necessarily be provided. Additionally, the efficiency of utilization of resources within a data packet network

used to transport circuit-switched data can also be comparatively low. In particular, trunking facilities are normally provisioned for peak utilization periods. This means that at any given time, some, or possibly most of the channels within a circuit-switched trunk may be idle. However, in order to facilitate recovery of the circuit-switched signals at an egress gateway, the entire trunk signal, including the idle channels, must be transported across the packet data network. This means that at least some of the packets (e.g. ATM cells, IP packets or Ethernet frames) carry null data.

United States Patent No. 5,936,965, which issued to Doshi et al. on August 10, 1999 teaches a system for supporting the transmission of multiple application-layer protocols through a single link using a single byte-stream. The multiple application-layer protocol types supported include asynchronous transfer mode (ATM) protocol data units (PDUs), synchronous transfer mode (STM) PDUs, and variable length (VL) PDUs, as well as subtypes included within these protocol types. Application-layer PDUs are processed at three intermediate protocol layers where the application layer PDUs are prepared, segmented, and repacked as asynchronous block multiplexing (ABM) PDUs. ABM PDUs include a type identification field. Cyclical redundancy checks and other error detection/correction techniques are optionally supported. ABM PDUs are multiplexed within a multiple protocol byte-stream. Multiple protocol byte-stream support is provided through a single link between a transmitter and receiver using a plurality of media, including coaxial cable, wireless,

optical fiber, hybrid fiber/coax, satellite, and twisted pair.

Thus the system taught by Doshi et al. is designed to operate in modern networks constructed using a data packet model and a layered architecture. Furthermore, Doshi et al. propose a proprietary ABM PDU format that is not supported by current network standards. Accordingly, the system of Doshi et al. can only operate on proprietary PHY-layer equipment, and thus is incapable of transporting data using the existing network infrastructure. Finally, Doshi et al. do not address the problem of adapting multiple legacy circuit-switched data streams for transport through a data packet network.

Accordingly, a method and apparatus that enables the protocol-independent extension of data services through a broadband packet network with efficient utilization of packet network resources, remains highly desirable.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a universal, i.e. protocol-independent, method and apparatus for enabling extension of data services through a broadband packet network.

Accordingly, an aspect of the invention provides a method of extending a data service through a broadband packet network. A data stream respecting the data service is received at an ingress gateway. The data stream is encapsulated within a container, which is then encapsulated within a protocol data unit (PDU) of the broadband packet

network. The PDU is then forwarded through the broadband packet network to an egress gateway.

5 A further aspect of the invention provides an apparatus for extending a data service through a broadband packet network. The apparatus includes: means for receiving a data stream respecting the data service at an ingress gateway; means for encapsulating the data stream within a container; means for encapsulating the container within a protocol data unit (PDU) of the broadband packet network; 10 and means for forwarding the PDU through the broadband packet network to an egress gateway.

Another aspect of the present invention provides an apparatus for extending a data service through a broadband packet network. The apparatus includes: means for receiving 15 sequential PDUs of the broadband packet network at an egress gateway from an ingress gateway; means for extracting a respective container from each received PDU; and means for reconstructing a data stream respecting the data service using the respective containers.

20 Another aspect of the present invention provides a system for extending a data service through a broadband packet network. The system includes an ingress gateway and egress gateway. The ingress gateway includes: means for receiving a data stream respecting the data service; means 25 for encapsulating the data stream within a container; and means for encapsulating the container within a protocol data unit (PDU) of the broadband packet network. Means are provided for conveying the PDU through the broadband packet network from the ingress gateway to the egress gateway. 30 The egress gateway includes: means for extracting a

respective container from each received PDU; and means for reconstructing the data stream using the respective containers.

5 In embodiments of the invention, encapsulation of the data stream comprises the steps of: accumulating a payload packet comprising a predetermined number of accumulated bytes of the data stream; and encapsulating the payload packet within the container.

10 The broadband packet network may be based on any one or more of the UDP/IP, TCP/IP, IP-MPLS, ATM, Ethernet and DOCSIS protocols, and the data stream may be formatted in accordance with any other communications protocol.

15 In some embodiments of the invention, a communications protocol of the data stream is known. In such cases, the predetermined number of bytes of the data stream forming each payload packet is preferably a function of the format of the data stream.

20 In embodiments in which the data stream is a known circuit-switched data stream comprising pulse code modulated PCM signals, the number of accumulated bytes forming each payload packet is determined by a number of channels and number of frames (or multi-frames or superframes) of the data stream. For example, In cases where the data stream is made up of Channel Associated Signaling, then the number of accumulated bytes forming each payload packet is preferably a function of the multi-frame count. In some embodiments, the number of accumulated bytes forming each payload packet may be equivalent to

$$P_s = (N_c \times N_m) \times n$$

Where: Ps = payload packet size;

Nc = number of channels;

Nm = number of frames (or superframes or multi-frames); and

5 n = an integer.

In other embodiments of the invention, the data stream is a known packet data stream comprising sequential PDU's of a packet network protocol. In such cases, the number of bytes forming each payload packet is preferably
10 an integer multiple of a number of bytes forming each PDU of the packet network protocol.

In some embodiments of the invention, a communications protocol of the data stream is unknown.

In embodiments of the invention, accumulation of a
15 payload packet may include: detecting an idle pattern; and when an idle pattern is detected, discarding bytes of the data stream corresponding to the detected idle pattern. The idle pattern may be known.

A known idle pattern may be embedded within the
20 data stream, in which case the step of detecting the idle pattern may include a step of monitoring successively received bytes of the data stream to detect the idle pattern. Alternatively, a known idle pattern may be a stimulus external to the data stream. The idle pattern may
25 be indicative of an idle channel within the data stream, and bytes within the indicated idle channel of the data stream can be discarded.

In embodiments of the invention, an idle notification is forwarded to the egress gateway. The idle notification preferably comprises information identifying detected idle patterns and corresponding idle channels. The idle notification may be forwarded within the container, or alternatively may be forwarded within a notification message independently of the container.

In embodiments of the invention, the idle pattern is unknown. In such cases, detection of an idle pattern preferably comprises monitoring each successive payload packet to detect a repeating pattern indicative of an idle condition of the data stream. The step of discarding bytes of the data stream may include discarding each successive payload packet in which the repeating pattern is detected. The steps of encapsulating payload packets within containers, encapsulating containers within PDUs and forwarding the PDUs to the egress gateway may be interrupted, and an idle notification sent to the egress gateway. Successive payload packets can then be monitored to detect the repeating pattern; and the steps of encapsulating payload packets within containers, encapsulating containers within PDUs and forwarding PDUs to the egress gateway can be resumed when the repeating pattern is no longer detected.

In embodiments of the invention, a sequence number is inserted into each successive container. At least one sequence number may be a reserved sequence number used to indicate a start of the data stream.

In embodiments of the invention, sequential PDUs of the broadband packet network are received at the egress

gateway from the ingress gateway. A respective container is extracted from each received PDU, and the data stream reconstructed using the respective containers.

Reconstruction of the data stream may comprise the

5 steps of: buffering each container in a jitter buffer; extracting a respective payload packet from each container; and appending each extracted payload packet to a payload packet of a previous container to reconstruct the data stream. The buffered containers may be sorted based on a

10 respective sequence number of each container. The respective sequence numbers of each buffered container may be monitored to detect a missing container. If a missing container is detected, a null payload packet may be appended to a previous payload packet of the reconstructed

15 data stream. The null payload packet may include AB-bits duplicated from the previous payload packet of the reconstructed data stream.

In embodiments of the invention, an inter-packet delay period is monitored between successively received

20 PDU's. A length of the jitter buffer can then be adjusted based on the inter-packet delay. The length of the jitter buffer is preferably adjusted during an idle period of the data stream.

In embodiments of the invention, reconstruction of

25 the data stream includes receiving an idle notification from the ingress gateway. The idle notification may include information identifying one or more of an idle indication and a corresponding idle channel of the data stream received by the ingress gateway, and the step of

30 reconstructing the data stream further comprises a step of

inserting null data into the corresponding idle channel of the reconstructed data stream following receipt of the idle indication. The null data may include the idle indication.

5 The idle notification may comprise an indication of an idle condition of the data stream received by the ingress gateway. In such cases, reconstructing the data stream may include any one or more of duplicating a last received payload packet, and inserting a provisioned idle pattern.

10 The idle notification may be received by the egress gateway encapsulated within a container, or alternatively may be received by the egress gateway within a notification message independently of a container.

15 Reconstruction of the data stream may resume based on containers extracted from received PDU's upon receipt of a container having a predetermined reserved sequence number.

20 An advantage of the present invention is that encapsulation of the data stream within containers renders the broadband packet network transparent to the source of the data stream (e.g. an enterprise site network), to thereby provide protocol-independent service extension through the broadband packet network. At the same time, idle suppression reduces congestion within the broadband
25 packet network by substantially eliminating the transport of null data.

BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended
5 drawings, in which:

FIG. 1. is a block diagram schematically illustrating an exemplary communications network within which the present invention may be deployed;

FIG. 2 is a block diagram schematically
10 illustrating encapsulation of a circuit-switched data stream for transport across the broadband packet network of FIG. 1;

FIG. 3 is a block diagram schematically illustrating reconstruction of a circuit-switched data
15 stream following transport through the broadband packet network of FIG. 1;

FIGs. 4a through 4c show respective exemplary formats of a container usable in the present invention;

FIG. 5 is a state diagram illustrating operation of
20 the ingress gateway in an embodiment of the invention in which idle suppression is not implemented;

FIG. 6 is a state diagram illustrating operations of the ingress gateway in an embodiment of the invention in which idle suppression is implemented and an idle pattern
25 is known; and

FIG. 7 is a state diagrams illustrating operations of the ingress gateway in an embodiment of the invention in

which idle suppression is implemented and an idle pattern is unknown.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention provides a method and apparatus which enables universal (i.e. protocol independent) extension of data services across a broadband packet network by transparently conveying data streams associated with such data services through the broadband packet network. FIG. 1 illustrates an exemplary communications network 2 within which the present invention may be deployed. In the embodiment illustrated in FIG. 1, a pair of legacy networks 4a,4b are connected to a broadband packet network 6 via respective gateways 8a,8b. The legacy networks 4a,4b may be, for example, respective domains of the PSTN, enterprise site networks, or physical connections to communications devices (e.g. a Plain Old Telephone Service [POTS] loop serving a telephone handset). Similarly, the legacy networks 4a,4b may operate in accordance with any legacy circuit-switched or packet-based communications protocol (e.g. El, T1, SNA, video, FR, ISDN etc.). Each gateway 8 is coupled to its respective legacy network 4a,4b via one or more physical interfaces 10a,10b, conforming to the connection standard applicable to the legacy network 4. The broadband packet network 6 may be any packet based network infrastructure including, wireless, optical fiber, co-axial cable hybrid fiber/coax,

or twisted pair, and using any of the packet based communication protocols (e.g. IP, ATM, Ethernet).

In practice, selection of the physical interface 10 between each gateway 8 and its respective legacy network 4 will be determined by provisioning, and thus the format of legacy data streams originating within these legacy networks 4 (e.g. unaligned, frame aligned, and multi-frame aligned) are known in advance. Similarly, the architecture and transport protocols of the broadband packet network 6 are determined by provisioning, and thus are known in advance. For the purposes of the present invention, it is assumed that issues of authentication, verification of service configuration, and data stream connectivity are handled by known methods. Thus, for example, for a trunked connection through the broadband packet network 6 between originating and destination points 12a,12b, it is assumed that known methods are utilized to set up a connection between the originating point 12a and a first (i.e. ingress) gateway 8a. It is also assumed that known methods are used to set up a connection between the destination point 12b and a second (i.e. egress) gateway 8b, and to establish signaling between the two gateways 8a,8b through the broadband packet network 6. It is further assumed that compatible versions of signal processing algorithms in accordance with the present invention are operative within each of the gateways 8a,8b so that signal adaptation can be performed properly in each gateway 8a,8b. Within this context, the apparatus of the present invention may be provided as a medium dependent adapter (MDA) 14 which may be coupled to the physical interface 10a, 10b of each

001230-501230

gateway 8a,8b so as to mediate traffic flow between the respective gateways 8a, 8b and legacy networks 4a,4b.

As will be described in greater detail below, the method of the present invention may be implemented either with, or without idle suppression. In embodiments of the invention in which idle suppression is implemented, it is possible to base the idle suppression functionality on either known (e.g., provisioned) idle patterns or discovered idle patterns, or both.

It will be appreciated that trunked data streams are typically bi-directional. However, in order to facilitate understanding of the present invention, and to simplify the present discussion, a first one of the gateways 8a,8b will be referred to as an ingress gateway 8a in which a received legacy data stream is packetized and launched into the broadband packet network. Similarly, a second one of the gateways 8a,8b is referred to as an egress gateway 8b in which received PDUs are processed to extract encapsulated payload packets which are assembled to regenerate the original legacy data stream.

In general, the present invention provides protocol independent adaptation services between legacy data streams and protocol data units (PDUs) of the broadband packet network. This is accomplished by treating a legacy data stream received at the ingress server 8a as a serial byte stream. As shown in FIG. 2, the serial byte stream 16 is split into successive packets 18, which may be of arbitrary size. The packets are inserted into respective containers 20. Each container 20 includes a header portion 22 that conveys data usable by the egress

gateway 8b for processing received containers 20 to regenerate the serial byte stream 16. In order to transport the containers 20 across the broadband packet network 6, each container is inserted into a respective
5 protocol data unit(PDU) 24 of the broadband packet network in a conventional manner.

As shown in FIG. 3, as each successive PDU 24 is received at the egress gateway 8b, the encapsulated container 20 is extracted from the PDU 24 and buffered
10 (e.g. in a jitter buffer, not shown). If necessary, the buffered containers 20 can be sorted (e.g. on the basis of a sequence number contained in the respective header portion 22 of each container 20). If desired, the inter-frame delay between arrival times of successive PDU's can
15 be monitored, and the length of the jitter buffer adjusted based on changes in the inter-frame delay. Finally, payload packets 18 encapsulated within each successive container 20 are extracted and processed as described below to produce a reconstructed serial byte stream 16b.

FIGs. 4a through 4c show respective exemplary formats for a container 20 usable in the present invention. As shown in FIG. 4a, each container 20 generally comprises a header portion 22 and a payload portion 26. The header portion 22 includes a sequence number field 28 and a
25 control field 30, each of which may be of a predetermined fixed length (e.g. 4 bits). Similarly, the payload portion 26 is preferably of a predetermined length. In general, the size of the payload portion 26 corresponds to the size of a payload packet 18, which may be arbitrary.
30 However, in embodiments in which legacy data streams are to be treated as frame aligned, or multi-frame aligned data,

the size of the payload portion 26 (and thus the payload packet 18) will preferably be based on the provisioned format of the respective interface 10. In particular, legacy signals 16 on a T1 interface 10 are divided into frames, and each frame is subdivided into N time slots. In that case, the payload portion 26 of the container 20 is preferably sized to carry an integer number of frames. Similarly, a multi-frame-aligned data stream is typically formatted as super-frames with each super-frame including a number of frames. If so, the size of the payload portion is preferably selected to accommodate an integer number of super-frames. This arrangement means that, if a container is lost during transport between the ingress and egress gateways 8a,8b, then an integer number of whole frames will be lost, so that problems resulting from sending incomplete frames into the receiving legacy network 4b are avoided. The relationship between frame alignment and payload packet size may be represented by the equation:

$$Ps = (Nc \times Nm) \times n$$

Where: Ps = payload packet size;

Nc = number of channels per frame;

Nm = number of frames (or super-frames); and

n = an integer.

Thus in general, the size of the payload portion 26 can be selected by provisioning on the basis of the physical interfaces 10 coupled to each of the ingress and egress gateways 8a,8b.

The sequence number field 28 of the container 20 is preferably used to carry a sequence number that is assigned to the container 20 in the ingress gateway 8a, and is used in the egress gateway 8b to ensure proper ordering of buffered containers 20 as well as detection of missing containers. The sequence number may be calculated by any of a variety of known means. For example, the ingress gateway 8a may be provided with a counter (not shown) that is incremented as each successive payload packet 18 is inserted into a respective container 20. The value of the counter can then be used as the sequence number relating to that payload packet 18 and thus inserted into the sequence number field 28 of the corresponding container 20.

As mentioned previously, within the egress gateway 8b, containers 20 are extracted from successively received PDUs 24 and buffered within a memory (not shown). The buffered containers 20 can then be sorted on the basis of their respective sequence numbers to ensure proper sequencing of payload packets 18 for assembly of the reconstructed legacy data stream 16b. Additionally, missing containers (not shown) (e.g. due to non-arrival of a PDU at the egress gateway 8b) can be readily detected by examination of the sequence numbers of buffered containers 20. At least one sequence number (e.g. "0000") may be reserved for specific signaling purposes. For example, a sequence number of "0000" may be inserted by the ingress gateway into the first container of a data stream, to thereby indicate (to the egress gateway 8b) the beginning of a flow of data. This facility can ensure proper reconstruction of the legacy data stream in the

egress gateway, even in the event that a notification message (indicating the start of a data flow) is lost.

The control code field 30 is used to convey information regarding the type of data stored in the payload portion 26, and therefore controls the processing of containers and reconstruction of the legacy data stream within the egress gateway. Exemplary control codes are shown in Table 1 below.

Control Code Value	Meaning
0000	payload includes data
0001	payload includes idle pattern indication
0010	unused
0011	payload includes nx64 idle indication bit map
0100	payload includes modem status leads
0101	payload includes fragmented frame
0110	payload includes end of fragmented frame
0111	payload includes multiplexing indication
1000	payload includes message notification
1001 - 1110	unused
1111	go to next byte for control byte

FIG. 4b illustrates a container which is formatted to include information 32 identifying idle channels within the legacy data stream 16, in addition to a payload packet 18. The idle channel information 32 may be, for

example, provided as a bit map (not shown) in which each bit corresponds to a channel of the legacy data stream 16. A value of "1" of a bit may, for example, indicate that the corresponding channel is idle, in which case a value of "0" would indicate that the corresponding channel is in service.

FIG. 4c illustrates a container usable in an embodiment where two or more legacy data streams 16 are multiplexed together at the ingress gateway 8a, and trunked through the broadband packet network 6 to the egress gateway 8b. In this case, each legacy data stream 16 is independently split into fixed length packets 18 which are encapsulated within containers 20 of the type illustrated in FIGs. 4a and 4b. These containers 20 are then encapsulated within a second container 20b which, in addition to a sequence number 28b and control code field 30b, also contains an interface field 34 and a channel number field 36 which enable the encapsulated containers 20 to be extracted and routed within the egress gateway 8b to thereby de-multiplex the packet streams prior to assembly of each of the respective reconstructed legacy data streams 16b.

As mentioned previously, the payload portion of each container 20 encapsulates a fixed length packet 18 of payload data, the size of which may be based on the provisioned format of the physical interface 10a connected to the ingress 8a. The process of packetizing legacy data streams 16 at the ingress gateway 8a is described below with reference to the state diagrams shown in FIGs. 5 through 7.

FIG. 5 is a state diagram illustrating operation of the ingress gateway for situations in which a received legacy data stream does not include any idle channels, and embodiments of the invention in which idle channel suppression is not implemented. In this case, successively received bytes of the incoming legacy data stream are accumulated (at S1) in a buffer (not shown) until a predetermined condition C1 is satisfied, at which time the accumulated bytes are encapsulated within a container 20, which is inserted into a PDU 24 for transport across the broadband packet network 6 (at S2). Upon completion of encapsulation of the accumulated bytes, the ingress gateway returns to accumulation of bytes of the incoming legacy data stream 16 within the buffer. The predetermined condition C1 will normally be based on the number of accumulated bytes stored in the buffer corresponding to the predetermined fixed size of each payload packet 18. In some embodiments, the buffer capacity may be provisioned to correspond to the predetermined fixed size of the payload packet 18, in which case the predetermined condition may be detection of the buffer being full. In other cases, the ingress gateway 8a may be provided with a counter (not shown) which operates to count successive incoming bytes (or frames, in the case of provisioned frame-aligned data signals), and the predetermined condition becomes satisfied when the counter reaches a predetermined threshold value.

FIGs. 6 and 7 are state diagrams illustrating operation of the ingress node 8a in embodiments of the invention that incorporate idle suppression. There are two fundamental reasons why idle suppression is important. The first is to save bandwidth in the broadband packet

network 6, and the other is to absorb variations in clocking at the ingress and egress gateways 8a,8b. Since some types of networks (e.g. IP networks) do not support cross-network clock synchronization, the respective gateway
5 clocks may vary. Idle suppression reduces traffic flow through the gateways 8a,8b, and thus allows time for queues in each of the gateways 8a,8b to re-center.

Idle suppression may be implemented in either of two ways: idle suppression where an idle pattern is known;
10 and idle suppression where the idle pattern is unknown. Each of these variations are described below in greater detail with reference to FIGs. 6 and 7 respectively.

FIG. 6 is a state diagram illustrating operations of the ingress gateway where an idle pattern is known. In
15 this case, the idle pattern (e.g. a flag indicating that a channel is idle) is known in advance (e.g. provisioned) and can be detected by the ingress gateway 8a within the received legacy data stream 16. For example, under the HDLC protocol, an idle channel is indicated by a known byte
20 (e.g. 0x7E or 0xFF) which is inserted into the corresponding channel, and never occurs in valid data (i.e. non-idle channels). Thus an idle channel within the data stream 16 can be detected by monitoring each successively received byte. Alternatively, an idle pattern may take the
25 form of a predetermined flag (or byte) that must be repeated a predetermined number of times to indicate an idle channel. For example, a Private Branch Exchange (PBX) may output a predefined byte when the channel is idle. If the ingress gateway 8a detects the byte during a predefined
30 period (or number of successively received frames) the channel may be considered to be idle.

As shown in FIG. 6, when an idle channel is detected at C2, the contents of that channel are dropped (i.e. discarded) at S3 so that only valid data are accumulated in the buffer. Optionally, a message can be sent to the egress gateway 8b (at S4) to indicate the idle channel and/or the idle pattern. This step enables the egress gateway 8b to insert the idle pattern into the buffered containers 20, so that the reconstructed data stream 16b transmitted by the egress gateway 8b fully reflects the contents of the data stream 16 received at the ingress gateway 8a.

The successive bytes received at the ingress gateway 8a can then be checked for the idle pattern (at S5), and, if no idle pattern is detected, the ingress node 8a continues (at S1) accumulating bytes of the data stream 16 in the buffer. As described above in respect of FIG. 5, when the predetermined condition C1 (e.g. buffer full or number of buffered bytes and/or frames) is satisfied, the accumulated data within the buffer is encapsulated into a container 20 for transmission through the broadband packet network 6 within a PDU 24.

For HDLC based data streams, the idle patterns (e.g. either 0x7E or 0xFF) can be removed from the data stream 16 by the ingress gateway 8a, so that only valid data is sent across the network 6. In addition, zero bit deletion/insertion and CRC16 or CRC32 checking can be performed at the ingress gateway 8a if desired. Since the HDLC service must understand the concept of a frame, and because HDLC frames can be quite large, there are two modes in which it can work. These are: Store-and-Forward, and Receive-and-Forward.

In a Store-and-Forward mode, an entire frame is received and validated in the ingress gateway 8a before being encapsulated and sent across the broadband packet network 6. This means that if an error is detected (e.g.,
5 invalid CRC) then the frame can be discarded without being encapsulated and sent across the broadband packet network 6.

In a Receive-and-Forward mode, bytes of a frame are encapsulated and forwarded across the broadband packet
10 network 6 as they are received. This means that if an error is detected at the ingress gateway 8a, a control message must be sent across the network 6 to inform the egress gateway 8b to abort the frame. The egress gateway 8b has the option of waiting until the entire frame
15 is received before transmitting, or it may begin part way through receiving the frame from the ingress side 8a.

FIG. 7 is a state diagram illustrating operations of the ingress gateway 8a where an idle pattern is unknown. In this case, the data stream 16 being received at the
20 ingress gateway 8a may be of an unknown type or may not have a predefined idle indicator (e.g. a video transmission or a bisynchronous signal), and thus an idle condition cannot be readily detected within the received data stream 16.

As shown in FIG. 7, successively received bytes are
25 accumulated (at S1) in a buffer (as described above) until the predetermined condition (C1) is satisfied. However, prior to encapsulating the buffered payload packet (S2), the ingress gateway 8a checks (at S6) for the presence of
30 repeating patterns within the buffered payload packet 18.

001250: 501250

5 This may be accomplished by ANDing the buffered payload packet 18 with a previous payload packet (not shown), and then analyzing the AND-operation result. If a repeating pattern is not detected (at C3), then the ingress gateway encapsulates and sends the buffered payload packet 18 as described above (state S7). On the other hand, if a repeating pattern is detected (at C4) in a predetermined number of successive payload packets 18, then the ingress gateway 8a determines that an idle condition exists. In this case, the ingress gateway 8a sends an Idle-condition notification message (at S8) to the egress gateway 8b and stops encapsulating payload packets and sending PDUs 24. In response to the Idle-condition notification message, the egress gateway 8b continues assembly of the reconstructed legacy data stream 16b by duplicating the last received payload packet 18. Following the sending of the Idle-condition notification message (S8), the ingress gateway 8a continues hunting for a "no idle" condition (S9), which may, for example, take the form of an absence of repeating patterns within successive payload packets 18. During this time idle probes (e.g. idle notification messages) may be sent to the egress gateway 8b. When the "no idle" condition is found (at C5), encapsulation and forwarding of payload packets resumes (at S1).

25 Using the above-described container formats and operations within the ingress and egress gateways 8a,8b, various system behaviors are possible, based primarily on the provisioned format of the data streams 16. Exemplary behaviors are described below, for each of unaligned, frame-aligned, and multi-frame-aligned data streams.

Unaligned data streams

Because an unaligned data stream is treated as a serial byte-stream, it is possible to trunk data streams across the broadband packet network 6 between gateways 8a,8b having differing interface formats. Thus the present invention is capable of performing format conversions of the data stream 16 as part of the trunking function. For example, a data stream 16 received at the ingress gateway 8a on any interface type (i.e. E1/T1/TTC2M/serial - where serial means V.35, V.11, V.24) can be reconstructed in the egress gateway 8b and transmitted through any other interface type. Only the order of the data is preserved during this conversion.

Provided that the line rates are identical at the ingress and egress gateway interfaces 10a,10b, it is possible to trunk data streams across the broadband packet network 6 between a T1/E1 interface and a serial interface. For example, a frame-aligned data stream having 6 timeslots received through a T1/E1 interface at the ingress gateway 8a can be carried across the broadband packet network 6 and reconstructed at the egress gateway 8b as a serial data stream, provided that the line rate of the serial interface is 384kbit/s.

In embodiments of the invention which include idle suppression, it is possible to trunk data streams between interfaces with different line rates, provided that the effective data rate through the broadband packet network 6 is less than or equal to the slowest line rate. For example, if a communications network 2 in which the ingress gateway 8a has a T1 interface 10a having a line rate of 512kbit/s, and the egress gateway 8b has a V.35

interface 10b having a line rate of 256kbit/s. If the effective data rate between the ingress and egress gateways 8a,8b is less than 256kbit/s, then it would be possible to trunk data across the broadband packet network 6 between these two interfaces 10a,10b.

For an unaligned data stream received through a serial interface, it is possible to carry modem status leads (i.e. RTS/DTR, etc) across the broadband packet network in special PDUs.

10 **Frame Aligned data Streams**

On T1 or E1 interfaces 10, legacy data streams 16 may be treated as frame or multi-frame aligned data streams. When frame aligned data streams are trunked across the broadband packet network 6 between T1 and E1 interfaces 10, the timeslot order as well as the data order is preserved, but multi-frame positioning may be lost. This means that data received at the ingress gateway 8a in timeslot 1 of frame 4 could be located in timeslot 1 of frame 8 of the reconstructed data stream 16b in the egress gateway 8b.

There is no requirement that the timeslot order be the same on each side. Thus data located in timeslots 1, 2 and 3 of the data stream 16a received at the ingress gateway 8a could be located in timeslots 17, 18 and 19 of the reconstructed data stream 16b in the egress gateway 8b. Similarly, there is no requirement that the timeslot order be contiguous on each side. Thus data located in timeslots 1, 2 and 3 of the data stream 16a received at the ingress gateway 8a could be located in timeslots 4, 18

and 23 of the reconstructed data stream 16b in the egress gateway 8b.

Multi-frame Aligned Data Streams

Multi-frame aligned data is very similar to frame aligned data streams, with the addition that the multi-frame positioning of the data is preserved. This means that data located in timeslot 4 of frame 6 of the data stream 16a received at the ingress gateway 8a is located in timeslot 4, frame 6 of the reconstructed data stream 16b in the egress gateway 8b.

There is no requirement that the timeslot order be the same on each side. Thus data located in timeslots 1, 2 and 3 of the data stream 16a received at the ingress gateway 8a could be located in timeslots 17, 18 and 19 of the reconstructed data stream 16b in the egress gateway 8b. Similarly, there is no requirement that the timeslot order is contiguous on either side. Thus data located in timeslots 1, 2 and 3 of the data stream 16a received at the ingress gateway 8a could be located in timeslots 4, 18 and 23 of the reconstructed data stream 16b in the egress gateway 8b.

Multi-frame aligned data streams can only be trunked between similar interfaces 10. For example, if the ingress gateway 8a receives a multi-frame aligned data stream through a T1 d4 interface 10, it may only trunk that data stream to an egress gateway 8b having a similar T1 d4 interface 10.

During processing of received PDUs 24, the egress gateway 8b operates to identify and extract the AB-bits from each super-frame encapsulated within the container 20.

If multiple super-frames are sent across the broadband packet network 6 in a single PDU, then the AB-bits should be extracted from each super-frame, in turn, and saved. These extracted AB-bits are stored for use in an event of frame loss. In particular, in a frame loss situation (i.e., a PDU 24 launched from the ingress gateway 8a fails to arrive at the egress gateway 8b), the stored AB-bits of the last-received super-frame can be inserted into the reconstructed data stream 16b so that multi-frame-alignment of the reconstructed data stream is preserved.

For multi-frame aligned data, the Channel Associated Signaling (CAS) bits may be carried across the broadband packet network 6 imbedded within containers 20 that are also carrying payload packets 18 of the data stream 16. For multi-frame aligned data it is important that the AB-bits be mapped correctly into the correct timeslots of frames 6, 12, 18 and 24 for T1 and timeslot 16 for E1.

The embodiment(s) of the invention described above
20 is(are) intended to be exemplary only. The scope of the
invention is therefore intended to be limited solely by the
scope of the appended claims.